

REMARKS

Claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 22-27, 29, 30, 32-34, 37, 38, and 40-56 are currently pending in the present application. Claims 1, 4-6, 9, 12, 14-16, 19, 22, 26, 29, 30, 32-34, 37, 38, 40-46, 50, 51, 53, and 55 have been amended. No claims are added or canceled by this amendment. Therefore, claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 22-27, 29, 30, 32-34, 37, 38, and 40-56 will remain pending in the application after entry of the foregoing claim amendments.

Claims 16, 22, 26, and 50 have been amended to correct typographical errors. Particularly, claim 16 has been amended to replace “1” with “12,” so that the claim now properly depends from claim 12. Claim 22 has been amended to replace the term “single log entry” with “single log”. Claims 26 and 50 have been amended to replace “allocator layer” with “allocation layer”.

Telephone Conversation With Examiner

Examiner Stace is thanked for the telephone conversation conducted on November 18, 2008. Proposed claim amendments were discussed. Cited art was discussed. It appears that the proposed claim amendments overcome the rejections based on the cited art.

Rejection under 35 U.S.C. § 103

Claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 29, 30, 32-34, 37, 38, 40-45, and 51-54 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 5,485,608 to Lomet et al. (“Lomet”) in view of “Efficient Locking for Concurrent Operations on B-Trees” (“Lehman”). Without conceding the merits of the rejection, Applicants have amended independent claims 1, 12, 22, 29, 38, 46, 51, 53, and 55 to further clarify the claimed subject matter. Support for the amendments to independent claims 1, 12, 22, 29, 38, 46, 51, 53, and 55 can be found, for example, at paragraphs [0100]-[0103] of the present application.

As noted in the Background section of the present application, the logging and recovery of a distributed data structure is important for ensuring reliability and scalability of the

underlying database system (*Specification* at ¶ [0015]). Distributed B-link tree data structures present unique challenges for logging processes when the logging occurs while the B-link tree is updated by a plurality of data transactions on different B-link tree nodes. For example, a single B-link tree operation may require multiple data transactions to implement. More particularly, for example, an insert operation for a B-link tree could not only cause a new data entry to be inserted, but could also lead to the split of one or more nodes along the path to the root (*id.* at ¶ [0100]). As a result, multiple data transactions are needed for certain B-link tree operations, and, therefore, multiple log records are required for logging all of the data transactions corresponding to the single B-link tree operation (*id.* at ¶ [0100]).

In a recovery mode, it is desirable to have the ability to determine whether a B-link tree operation corresponding to multiple data transactions has been fully completed. If all data transactions corresponding to the single operation are contained in the log, then the operation fully completed prior to system failure. However, if all of the data transactions corresponding to the single operation are not contained in the log, the operation did not fully complete, and, therefore, a logical undo of the operation must be performed to properly recover the B-link tree data structure (*id.* at ¶¶ [0101]-[0103]). To solve this problem and others, and to overcome other deficiencies associated with existing logging techniques with respect to B-link trees, the claimed subject matter implements logging methods and systems that associate log entries corresponding to a single B-link tree operation for determining whether the single B-link tree operation has been completed.

For example, independent claims 1, 12, 29, 38, 51, and 53 each recite, in part, *(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single B-link tree operation, and (2) associating the log entries with each other for use in determining whether the single B-link tree operation has been completed.* For example, a B-link insert operation requires multiple log records (*id.* at ¶ [0103]). A prefix can associate these log records, such that if a prefix of these log records is found, the insert operation can be logically completed (*id.*). But when it is known that the operation did not complete, as

evidenced by a partial prefix in the log, a logical undo of the insert operation can be performed (*Id.*).

Lomet does not teach ***(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single B-link tree operation, and (2) associating the log entries with each other for use in determining whether the single B-link tree operation has been completed.*** Referring to the Summary of the Invention section of Lomet, Lomet teaches recording, in a log, state identifiers of persistent storage sections (Lomet at col. 2, lines 56-65). Each log can be used for recovery independently from the other logs by applying to each of the sections in memory only the actions from that log which have the same state identifier value as is stored in the most recent version of that section (Lomet at col. 2, line 66, to col. 3, line 6). There is no teaching or suggestion in Lomet of the claimed feature of generating log entries corresponding to data transactions associated with a single B-link tree operation. Although Lomet teaches the use of logging with respect to sections of data in general, the use of B-link trees provides unique challenges to logging and recovery, which are addressed by the claimed subject matter and are neither taught nor suggested by Lomet.

More particularly, the claimed subject matter includes the feature of associating the log entries with each other for use in determining whether a single B-link tree operation has been completed. This claimed feature is important for the reasons set forth above for proper recovery with respect to B-link tree operations, such as insert operations, requiring more than one data transaction. There is no teaching or suggestion in Lomet of such a feature.

Lehman does not overcome the shortcomings of Lomet. Lehman generally teaches various B-tree data structures. However, there is no teaching or suggestion in Lehman of the claimed features of ***(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single B-link tree operation, and (2) associating the log entries with each other for use in determining whether the single B-link tree operation has been completed.*** Therefore, for at least the foregoing reasons, even if combined, Lomet and Lehman do not teach or suggest the features of independent claims 1, 12, 29, 38, 51, and 53.

Accordingly, for the foregoing reasons, Applicants respectfully request the withdrawal of the rejection of claims 1, 12, 29, 38, 51, and 53 under 35 U.S.C. § 103(a).

Claims 2, 4-6, 9, 10, 14-16, 19, 20, 30, 32-34, 37, 40-45, 52, and 54 depend upon one of claims 1, 12, 29, 38, 51, and 53. Therefore, for at least the reasons set forth for claims 1, 12, 29, 38, 51, and 53, Applicants respectfully submit that even if combined, Lomet and Lehman do not teach or suggest the features of dependent claims 2, 4-6, 9, 10, 14-16, 19, 20, 30, 32-34, 37, 40-45, 52, and 54. Accordingly, for the foregoing reasons, Applicants respectfully request the withdrawal of the rejection of claims 2, 4-6, 9, 10, 14-16, 19, 20, 30, 32-34, 37, 40-45, 52, and 54 under 35 U.S.C. § 103(a).

Claims 22, 23, 26, 46, 47, 50, 55, and 56 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Lomet in view of U.S. Patent No. 5,434,994 to Shaheen et al. ("Shaheen").

Claims 22, 46, and 55 each recite, in part, *(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single operation on a data structure, and (2) associating the log entries with each other for use in determining whether the single operation on the data structure has been completed.* Applicants respectfully submit Lomet does not teach these claimed features. Referring to the Summary of the Invention section of Lomet, Lomet teaches recording state identifiers of persistent storage sections in a log (Lomet at col. 2, lines 56-65). Each log can be used for recovery independently from the other logs by applying to each of the sections in memory only the actions from that log which have the same state identifier value as is stored in the most recent version of that section (Lomet at col. 2, line 66, to col. 3, line 6). For reasons similar to those set forth above with respect to independent claims 1, 12, 29, 38, 51, and 53, Applicants respectfully submit that there is no teaching or suggestion in Lomet of the claims 22, 46, and 55 feature of generating log entries corresponding to data transactions associated with a single operation on a data structure.

Shaheen does not overcome the shortcomings of Lomet. Shaheen teaches merging logs into different servers (Shaheen at col. 4, lines 62-66, and col. 7, lines 50-55). However, there is no teaching or suggestion in Lehman of the claimed features of *(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single operation on*

a data structure, and (2) associating the log entries with each other for use in determining whether the single operation on the data structure has been completed.

In the telephone conversation of November 18, 2008, Examiner Stace directed Applicants' attention to column 7, lines 50-55, of Shaheen. This portion of Shaheen refers to the merging of logs from subordinates by a coordinator server, and the distribution of the merged log to all subordinates (*id.* at col. 7, lines 6-21 and 49-52). Each subordinate performs updates missing from its replica (*id.* at col. 7, lines 51 and 52). A single transaction is used to execute the missing updates (*id.* at col. 7, lines 52-54). Thus, this portion of Shaheen is directed to merging logs from several subordinate servers. There is no disclosure or suggestion in this portion of Shaheen, or any other portion of Shaheen, of *generating a plurality of log entries corresponding to a plurality of data transactions associated with a single operation on a data structure* in accordance with the claimed subject matter. In addition, there is no disclosure or suggestion in this portion of Shaheen, or any other portion of Shaheen, of *associating the log entries with each other for use in determining whether the single operation on the data structure has been completed* in accordance with the claimed subject matter.

Therefore, for at least the foregoing reasons, even if combined, Lomet and Shaheen do not teach or suggest the features of independent claims 22, 46, and 55.

Accordingly, for the foregoing reasons, Applicants respectfully request the withdrawal of the rejection of claims 22, 46, and 55 under 35 U.S.C. § 103(a).

Claims 23, 26, 27, 47, 50, and 56 depend upon one of claims 22, 46, and 55. Therefore, for at least the reasons set forth for claims 22, 46, and 55, Applicants respectfully submit that even if combined, Lomet and Shaheen do not teach or suggest the features of dependent claims 23, 26, 27, 47, 50, and 56. Accordingly, for the foregoing reasons, Applicants respectfully request the withdrawal of the rejection of claims 23, 26, 27, 47, 50, and 56 under 35 U.S.C. § 103(a).

Claims 24, 25, 48, and 49 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Lomet in view of Shaheen, further in view of Lehman. Claims 24, 25, 48, and 49 depend upon one of claims 22 and 46. For at least the reasons set forth above with respect to claims 22

and 46, Applicants respectfully submit that Lomet and Shaheen, either alone or in combination, do not teach or suggest the features of dependent claims 24, 25, 48, and 49. Particularly, Lomet and Shaheen do not teach or suggest the claimed features of *(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single operation on a data structure, and (2) associating the log entries with each other for use in determining whether the single operation on the data structure has been completed.*

Lehman does not overcome the shortcomings of Lomet and Shaheen. Lehman generally teaches various B-tree data structures. However, there is no teaching or suggestion in Lehman of *(1) generating a plurality of log entries corresponding to a plurality of data transactions associated with a single operation on a data structure, and (2) associating the log entries with each other for use in determining whether the single operation on the data structure has been completed.* Therefore, for at least the foregoing reasons, even if combined, Lomet, Shaheen, and Lehman do not teach or suggest the features of claims 24, 25, 48, and 49.

Accordingly, for the foregoing reasons, Applicants respectfully request the withdrawal of the rejection of claims 24, 25, 48, and 49 under 35 U.S.C. § 103(a).

DOCKET NO.: MSFT-2557/304882.01
Application No.: 10/674,676
Office Action Dated: September 15, 2008

PATENT

CONCLUSION

In view of the foregoing, Applicants respectfully submit that the claims are allowable and that the present application is in condition for allowance. Reconsideration of the application and an early Notice of Allowance are respectfully requested. In the event that the Examiner cannot allow the present application for any reason, the Examiner is encouraged to contact the undersigned to discuss the resolution of any remaining issues.

Date: December 15, 2008

/Joseph F. Oriti/
Joseph F. Oriti
Registration No. 47,835

Woodcock Washburn LLP
Cira Centre
2929 Arch Street, 12th Floor
Philadelphia, PA 19104-2891
Telephone: (215) 568-3100
Facsimile: (215) 568-3439